Numerical Range Facets Partition: Evaluation Metric and Methods

Xueqing Liu, Chengxiang Zhai
University of Illinois at Urbana-Champaign
Urbana, IL 61801
xliu93,czhai@illinois.edu

Wei Han, Onur Gungor
WalmartLabs
Sunnyvale, CA 94806
whan,ogungor@walmartlabs.com

arXiv:1610.10000v1 [cs.IR] 31 Oct 2016

ABSTRACT

Faceted browsing is a very useful interface component provided in many of today's search engines. It is especially useful when the user has an exploratory information need or have a clear preference for certain attribute values. Existing work has tried to optimize faceted browsing systems in many aspects, but little work has been done on optimizing the partitions of numerical facet ranges (e.g., price ranges of products). In this paper, we introduce the new problem of numerical facet range partition and formally frame it as an optimization problem. To enable quantitative evaluation and reuse of search log data, we propose an evaluation metric based on user's browsing cost when using the suggested ranges for navigation. We further propose and study multiple methods to computationally optimize the partition by leveraging search logs. Experimental results on a two-month search log from a major e-Commerce engine show that learning can significantly improve the performance over baseline.

CCS Concepts

•**Information systems** → **Search interfaces;** Query suggestion;

Keywords

Faceted browsing; Query log analysis; Information retrieval models; Non-smooth optimization; Learning to rank

1. INTRODUCTION

Querying and browsing are two complementary ways of information access. As a convenient way to support browsing, faceted browsing has now become a critical part in today's search engines. Faceted browsing is especially useful in exploratory search, where the user would issue very simple queries in order to examine more entities and make the best decision. In these cases, it is difficult for the ranking itself to guess the user's actual information need. With the help of faceted browsing, user can then conveniently specify her need by selecting from suggested facets. Faceted browsing has also become increasingly popular due to the rise of small screen devices, where typing queries is less convenient than clicking on facets. A standard faceted browsing system consists of a list of facets, including categorical facets and numerical facets. Each facet is followed by its facet values, and user can navigate by selecting one or more facet values from the listed ones. An example of a faceted browsing system is shown in Figure 1.

Figure 1: A faceted browsing system, picture from [5]

There has been a significantly amount of work on optimizing faceted browsing systems [7, 8, 11, 14]. However, most of these works focus on optimizing categorical facets rather than numerical facets. As is shown in Figure 1, a numerical facet contains a list of non-overlapping ranges, partitioning the ranked list into k refined sub-lists. Therefore, the optimization problem is equal to finding the $k-1$ best separating values. Apparently, existing methods on categorical facets are not designed for solving this problem. As a result, numerical range partition is a new research problem which has not been solved by any existing work.

Despite the fact that numerical facet is one single component of the entire faceted browsing system, it has been playing an indispensable role in many applications. Examples of numerical facets include price of products, timestamps of events, distance between two geolocations, and any other artificially defined quantity (e.g., popularity of topics). Therefore, numerical range partition is also a general problem with wide range of application scenarios.

In this paper, we conduct the first systematic study on numerical range partition. Since this problem is completely new, we need to first solve the problem of evaluation. Generally speaking, IR evaluations are of two folds: online and offline. Online evaluation usually compares an algorithm against baseline using controlled experiments such as A/B testing. Although this approach is successful, it is unfortunately time-consuming and expensive. On the other hand, offline approach would evaluate an algorithm by defining metrics (e.g., MAP, NDCG) on previous search logs and reusing previously labelled relevance judgement. Since of-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WOODSTOCK '97 El Paso, Texas USA

© 2016 ACM. ISBN 123-4567-24-567/08/06…$15.00

DOI: 10.475/123_4

fline approach is inexpensive and tremendously faster than online approach, it has been used by the majority of previous work. We thus choose the offline approach to evaluate our partition problem.

However, there exists one challenge in offline evaluation: different from ranking problem, it is not straightforward how to obtain reusable relevance judgements in our problem. Indeed, numerical ranges can have infinite number of possible values, and it is impossible to enumerate all the (query, ranges) pairs. We can only try to guess user's response when she sees a set of numerical ranges. We notice that two assumptions can help us make this guess. First, user should have some brief idea about the numerical range she is more interested in (e.g., laptops below 200$). Second, user should try to minimize the time she spends before reaching the first relevant entity. Using these assumptions, we define our evaluation metric as the minimum number of entities user has to browse before reaching the first relevant entity when using this partition.

Once the evaluation metric is settled, we explore two methods for generating the partition. We try to use the two methods to answer one research question: given access to previous logs, i.e., training data, can we improve our evaluation metric on testing data? Specifically, in the first approach, we generate the partition that optimizes the expectation of our evaluation metric, where the probability values in expectation are estimated from the training data; in the second approach, we introduce some parameters and learn these parameters from training data. We conduct experimental study on a two-month search log from Walmart.com[1]. By comparing against a simple baseline without training but with reasonable performance, we conclude that leveraging training data can indeed improve the evaluation metric on test collection.

The contributions of this paper are of three folds:

- We propose to first study the new problem of numerical facet range partition;
- We define the first offline evaluation metric that quantifies the performance of a numerical facet partition algorithm on search logs;
- We propose two methods for suggesting numerical facet ranges. Experimental results against baseline show that they can achieve significantly better performance on our proposed metric;

2. RELATED WORK

There has been a significant amount of work on optimizing the faceted browsing interface [7, 11, 8]. However, all of these works assume the facet values (including numerical facets) are pre-defined and fixed, and the optimizations are either on ranking facets [11], or ranking (facet, value) pairs [7], or selecting a fixed number of facets to display [8]. To the best of our knowledge, no previous work has studied optimization of numerical range partition. The results generated by our work can be applied on top of any of these systems [7, 11, 8]. That is to say, our work is along a different direction with these previous works, thus should not be compared.

Our proposed evaluation metric employs the principle of minimizing user's browsing cost. This principle has served as an important strategy in many interface optimization problems [11, 7, 15]. For example, Liberman et al. [8] define the expected browsing cost of a facet as the rank of relevant document in the refined list. In contrast, Kashyap et al. [7] define it as the total number of entities in the refined list. Between the two approaches, we believe the former one more precisely defines the browsing cost, therefore we follow the former approach to define our own evaluation metric, although the latter serves as our learning objective function.

Our problem is remotely related to generating histograms for query optimization [6, 1, 9]. Histograms are used for data compression, therefore its evaluation metric is completely different from ours. Among these works, Jagadish et al. [6] use dynamic programming, i.e., the same optimization method we use for our first method. Some recent work [1] leverages an approximation technique to obtain linear running time, avoiding quadratic time cost in DP. However, this technique is not applicable to our problem due to the inherent difference in objective function. As a result, we have not achieved a lower time cost for our first method.

3. PROBLEM FORMULATION

In this section, we formally define the numerical range partition problem and introduce notations that we will use throughout the rest of the paper.

3.1 Input and Output

Suppose we have a working set of entities $E = \{e_1, \cdots, e_{|E|}\}$ that user would like to query on. At time i the user issues a query q^i, the search component[2] returns a list of retrieved entities $E^i \subset E$. Our problem is: for one specific numerical facet (e.g., price, customer rating), output a partition that will be displayed under the specific numerical facet as in Figure 1. We will assume from now on that the specific facet is fixed, and refer to it using terminology "the numerical facet value of entity e".

To formalize this problem, we further require that the total number of ranges is given as input. This is for fair comparison on the output partition. Indeed, it would be unfair to compare [0, 200), [200, 400) against [0, 100), [100, 200), [200, 300), [300, 400]. The latter is almost certainly better, because user can always use it to zoom into a better refined list than the former one. In real world applications, the number of ranges can either be given by the search engine, or defined by the user.

In summary, our input and output are: **Input**: (1) query q^i; (2) a fixed ranking algorithm, and ranked list E^i it retrieves for q^i; the numerical value of each entity $e \in E^i$, denoted as $v(e)$; the rank of each entity $e \in E^i$, denoted as $rank(e)$. In addition, we may also have access to the relevance scores, denoted as $p(e)$, i.e., the probability for entity e to be relevant[3]; (3) the number of ranges k. **Output**: k ranges, denoted by $k-1$ increasing separating values $P^i = [p_1, \cdots, p_{k-1}] \in \mathbb{R}^{k-1}$.

3.2 Reusing Query Logs

As discussed in Section 1, the problem of numerical facets partition has not been studied before, therefore we need to first answer the evaluation question: how can we quantitatively measure the quality of a partition?

[1] www.walmart.com

[2] The entity ranking function does not rely on the numerical facet partition, therefore we assume it is pre-defined and fixed.

[3] Notice the difference between notation p_j and $p(e)$. We use p_j with subscript to denote partition values, and $p(e)$ to denote the entity relevance models.

We start answering this question by looking into traditional IR metrics in ranking problem. In traditional IR, the relevance of each entity is labelled against a query, and this relevance judgement is reused on a test collection to evaluate the performance of a new ranking algorithm. However, in our problem, it is not straightforward how to obtain reusable relevance judgements. This is because there are infinite number of possible numerical ranges, thus it is impossible to enumerate all the (query, ranges) pairs. The best we can do is to guess which range the user would select if we show her a set of ranges.

Fortunately, we observe two assumptions on user behavior that can help us make this guess. First, upon seeing all the ranges, user should have a brief idea about which one she is more interested in, therefore, she should be able to recognize the unique range that contains her relevant entity and select it. Second, user should always try to minimize the time she spends before reaching the first relevant entity, therefore after selecting the numerical range, she would sequentially browse the refined list until reaching the relevant entity. To simplify our problem, we focus on evaluating the browsing cost for the first relevant entity, although our metric can be generalized to the case with more than one relevant entities.

The above assumptions may not be always true in reality, yet they are reasonable to make when our main purpose is to perform comparative studies between different partitioning algorithms. This is because if there is any bias introduced through these assumptions, the bias is unlikely favoring any particular algorithm, making the relative comparison still meaningful. This is essentially the same argument as the Cranfield evaluation methodology for information retrieval [12], where we obtain a reusable test collection for comparing different algorithms at the cost of making a few unrealistic assumptions about users and relevance judgement.

3.3 Quantifying the Evaluation Metric

Based on the assumptions defined in Section 3.2, suppose we have access to a user query log with query $\mathcal{Q} = \{q^1, \cdots, q^n\}$ and entity lists $\mathcal{E} = \{E^1, \cdots, E^n\}$. In addition, suppose the first relevant entity under each query is e^i. We want to evaluate an algorithm which generates partitions $\mathcal{P} = \{P^1, \cdots, P^n\}$. Given query q^i, the algorithm displays $P^i = [p_1, \cdots, p_{k-1}]$. The user is able to identify the unique range $[p_{j-1}, p_j)$ containing $v(e^i)$, i.e., the relevant entity's numerical facet value. After selecting $[p_{j-1}, p_j)$, E^i is refined to the sub-list containing only entities $e \in E^i$ with numerical value $v(e) \in [p_{j-1}, p_j)$. Since the user would sequentially browse the refined list until finding e^i, we define her browsing cost as the rank of e^i in the refined list. We name this value the *refined rank* for query i, denoted as RR_i:

$$RR_i \quad = \quad Rank\ of\ e^i\ in\ its\ refined\ list$$

We further obtain the *average refined rank* ARR:

$$ARR \quad = \quad \frac{1}{n}\sum_{i=1}^{n} RR_i \tag{1}$$

ARR will serve as the main evaluation metric we use throughout the rest of this paper.

4. METHODS

One straightforward method for generating a partition is to make each refined list contain equal number of entities. This approach is also named *equi-depth bining* or *quantile* method in literature [9]. Despite its simpleness, the quantile method actually perform reasonably well given no extra knowledge on user preference or training data. Consider the example in Table 1. For each of the two queries "refurbished laptop" and "gaming laptop", search engine displays the same partition [500, 600, 700, 800, 1,000]. However, users looking for refurbished laptops would presumably be more interested in price under 500$. When using this partition, they would often have to scan through a list containing more than 1,000 entities. In contrast, the quantile method would partition the two retrieved lists into ranges containing 321 and 408 entities, respectively. Therefore, quantile method is already better than the currently approach in the search engine.

Range	# retrieved entities	
	refurbished	*gaming*
Under 500$	1,426	306
500 - 600$	111	120
600 - 700$	75	95
700 - 800$	92	101
800 - 1,000$	90	275
Above 1,000$	234	1,949

Table 1: A snapshot from Amazon.com on 2016-01-21. The two queries are "refurbished laptop" and "gaming laptop". Column 1 shows the same numerical ranges the search engine displays for the two queries, column 2 and column 3 show the respective numbers of refined entities inside each range.

However, when we do have access to user preferences and training data, can we leverage them to improve the performance against quantile method? In this section, we propose two methods to study this problem.

4.1 First Method: Optimizing Expected ARR

Once the evaluation metric is defined as ARR (Equation 1), we should try to minimize the ARR in the test collection. Imagine if the relevant entity e^i was known, this means we should partition the retrieved list so the refined sub-list containing e^i only contains e^i itself. In other words, one of the ranges is $[v(e^i)-\epsilon, v(e^i)+\epsilon)$ where $\epsilon > 0$ is a small enough value. ARR in this imaginary scenario is equal to 1. But since the relevant entity is not known before the partition, what strategy should we use instead?

Suppose, however, we have access to the estimated probability each entity $e \in E^i$ is relevant, i.e., $p(e)$ and $\sum_{e \in E^i} p(e) = 1$. Under the simple assumption that there is only one relevant entity (Section 3.2), the expectation of RR_i when using partition $P^i = [p_1, \cdots, p_{k-1}]$ is thus:

$$\mathbb{E}_{P^i}[RR_i] \quad = \quad \sum_{e \in E^i} p(e) \times (Rank\ of\ e\ in\ its\ refined\ list)$$

We propose our first method: return P^i that minimizes $\mathbb{E}_{P^i}[RR_i]$. To minimize $\mathbb{E}_{P^i}[RR_i]$, notice that although the domain $P^i \in \mathbb{R}^{k-1}$ is continuous, the problem itself is discrete. For example, suppose there are three retrieved entities: position 1=100$, position 2=200$, position 3=300$, $k = 2$, and $p(1) = 0.4, p(2) = 0.3, p(3) = 0.3$. Notice that $\mathbb{E}_{P^i}[RR_i]$ stays the same as p_1 increases from 200.01 to 299.99. In fact, we only need to search for two points in the domain, i.e., $p_1 = 150$ and $p_1 = 250$. Therefore, this is a combinatorial optimization problem that searches for the best $\mathbb{E}_{P^i}[RR_i]$ among all the sets of size $k-1$ out of $|E^i|-1$ candidates.

We can use brute-force search for the combinatorial opti-

mization, however, the time complexity would be $O(\binom{|E^i|-1}{k-1} + |E^i|^3 \log |E^i|)$, where the $|E^i|^3 \log |E^i|$ is for sorting and obtaining the ranks of each entity in each possible range. When $|E^i|$ is large, this time cost is thus undesirable. However, since the objective function $\mathbb{E}_{P^i}[RR_i]$ can be written as the sum of k expected costs, where the k-th cost is independent from previous $k-1$ costs, we can apply dynamic programming here. The time cost for dynamic programming is reduced to $O(k|E^i|^2 + |E^i|^3 \log |E^i|)$.

One may wonder why we do not use greedy algorithm here. There are two reasons: first, greedy algorithm generally leads to sub-optimal solutions[4]; second, the computational cost of greedy algorithm is $O(k|E^i| + |E^i|^3 \log |E^i|)$, which remains large since it still has to compute ranks of each entity in each possible range.

The only question left is how to estimate $p(e)$. Basically, we can estimate $p(e)$ by applying any previous work in probabilistic retrieval models (language models) [17], as well as previous search logs. We will explore multiple ways for estimating $p(e)$ in Section 5.

4.2 Parameterizing the Problem

Given the definition of ARR (Equation (1)) and access to training data, another idea is: can we model the ARR using a set of parameters, and learn the parameters from training data, in a similar way as *learning to rank* (e.g., SVM-map [16])? At first sight, there does not seem to exist a very straightforward solution to the parameterization. One may think it is possible to define constant separating values $[p_1, \cdots, p_{k-1}]$ as the parameters. However, we can see that this is a sub-optimal strategy by considering the example in Table 1 again. Since users looking for refurbished laptops would be more interested in price under 500$, they would be benefitted from better defined ranges below 500$ than above 1,000$. For the same reason, users looking for gaming laptops would hope ranges above 1,000$ is better defined than under 500$. Therefore, it is impossible for the same partition $[p_1, \cdots, p_{k-1}]$ to satisfy both needs at the same time.

As discussed at the beginning of Section 4, the simple quantile method instead have reasonable performance at adapting to different information needs. Under assumptions in Section 3, quantile method would avoid user to scan over 1,000 entities in the worst case. More importantly, the quantile method reveals an alternative way to parameterize the problem: to separate the retrieved list by the relative *percentage* instead of absolute values[5]. The percentage parameters of the quantile method are $[1/k, \cdots, (k-1)/k]$, which corresponds to one point in the simplex space Δ^k. In general, we can define the percentages to be any point $R = \{r_1, \cdots, r_{k-1}\}$ (where $0 < r_1 < \cdots < r_{k-1} < 1$) in the simplex space. One would then come up with a few questions: does the quantile method $[1/k, \cdots, (k-1)/k]$ already achieve the optimal ARR among all $R \in \Delta^{k-1}$? If not, can we achieve better ARR than quantile by optimizing ARR (1) on training data?

4.2.1 Optimizing with Respect to R

Similar as most of the ranking metric (e.g.,NDCG, MAP, etc.), it is also difficult to directly optimize ARR, because ARR is a non-smooth objective function with respect to R. Indeed, if the relevant entity is near the boundary defined by R, and we change R with a small enough value, the relevant entity would jump from one range to another. For example, suppose there is only one query with three retrieved entities: position 1=100$, position 2=200$, position 3=300$; the relevant entity is position 2 and $k = 2$. If we change $R = [0.66]$ to $R = [0.67]$, the partition switch from {{1}, {2,3}} to {{1,2}, {3}}, and ARR switch from 1 to 2.

Non-smooth optimization. To solve the non-smooth optimization problem, first notice that ARR can be non-smooth everywhere, instead of being non-smooth at only a few points[6]. There exist a few derivative-free algorithms for solving the optimization problem in this case. Two important algorithms are the Powell's conjugate direction method [3] and the Nelder-Mead simplex method [10].

Time complexity to directly optimize ARR. Although the non-smooth optimization problem is solved by the aforementioned algorithms, we will see that the time complexity for these algorithms to directly optimize ARR is still very large. This time complexity can be written as $O(T_1 T_2)$, where T_2 is the average time cost to compute ARR on a specific point, and T_1 is total number of times this operation (computing ARR on a specific point) is needed. In other words, T_1 is the actual time complexity of the algorithm. Meanwhile, we can observe from Equation (1) that $T_2 = O(nm)$, where m is the average number of retrieved entities $|E^i|$ for each query q^i. That is, we must recompute the entire ARR from scratch at every point. The reason behind it is: on one hand, whenever we are at a new point, every RR_i in Equation (1) could have changed (as we discussed above, a small enough change in R can lead to a significant change in RR_i and ARR), so we have to recompute the RR_i in every single query; on the other hand, every such recomputation takes $O(m)$, because we need to recompute the ranges based on relative percentages R instead of absolute values (this is also explained in Section 4.4, $O(m)$ is the time complexity of k-th smallest element algorithm). As a result, the time complexity for any optimization algorithm to directly optimize ARR is $O(T_1 mn)$. In real world search engine, both m and n could be very large, this time cost is thus undesirable.

4.2.2 Optimizing the Surrogate Objective Function

As discussed in Section 4.2.1, the algorithm for directly optimizing ARR takes $O(T_1 mn)$. In order to reduce this cost, we propose a three-step process which turns ARR into a surrogate objective function with significantly reduced time cost.

Step 1: Normalization. For each query i, we normalize the rank of relevant entity in Equation (1) by the total number of retrieved entities E^i, denoted as $\overline{RR}_i$:

$$\overline{RR}_i = \frac{Rank\ of\ e^i\ in\ its\ refined\ list}{|E^i|}$$

Step 2: Upper bound. Notice the rank of the relevant entity in its refined list is upper bounded by the total number

[4]An example: position 1=400$, 2=100$, 3=200$, 4=300$, $p(1) = p(2) = 0.2, p(3) = p(4) = 0.3, k = 3$. $\mathbb{E}_{P^i}[RR_i]$ of optimal solution is 1.2 but greedy is 1.3.

[5]We slightly abuse the term *percentage*, which originally only refers to the numerator of a fraction when the denominator is 100, but here we use it to refer to any fractional values between 0 and 1.

[6]Therefore our optimization cannot be solved in the same as Lasso [13] which uses sub-gradient descent.

of entities in that range:

$$\overline{RR}_i \quad \leq \quad \frac{\#\ entities\ in\ e^{i\prime}s\ refined\ list}{|E^i|} \qquad (2)$$

Step 3: Limit approaching infinity. Notice that as $|E^i|$ goes to infinity, the limit of R.H.S. of (2) is $r_j - r_{j-1}$, where $[r_{j-1}, r_j)$ is the specific range in R that contains the relevant entity. Denote z^i as the relative percentage of $v(e^i)$ in the retrieved list[7], we can approximate R.H.S. of (2) as:

$$C^i(R) = \sum_{j=1}^{k} \mathbb{1}[r_{j-1} \leq z^i \leq r_j] \times (r_j - r_{j-1})$$

We now obtain the surrogate objective function:

$$\begin{aligned} C_n(R) &= \frac{1}{n}\sum_{i=1}^{n} C^i(R) \\ &= \sum_{j=1}^{k} \Delta r_j \times (F_n(r_j) - F_n(r_{j-1})) \qquad (3) \end{aligned}$$

Where $\Delta_j = r_j - r_{j-1}$, and $F_n(r_j) = \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}[z^i < r_j]$ is the empirical conditional distribution function (CDF) of z^i. Compared with the $O(T_1 mn)$ time cost for directly optimizing ARR (Section 4.2.1), the time cost for optimizing $C_n(R)$ is reduced to $O(nm + T_1 k \log n)$. This is because we can first compute and cache the empirical CDF. After that, to compute $C_n(R)$ at any new point R, we only need a binary search on the cached CDF.

4.2.3 Bounds on $C_n(R)$

The Dvoretzky-Kiefer-Wolfowitz inequality [4] bounds the probability that the empirical CDF F_n differs from the true distribution F. Following the DKW inequality, we are able to prove a few bounds on $C_n(R)$. These bounds provide useful insights on the convergence rate and sample complexity of $C_n(R)$ on large scale datasets, and are shown in Theorem 1 and Theorem 2 in Appendix.

4.3 Learning to Partition with Decision Tree

In Section 4.2 we propose to parameterize the problem under the same percentage parameter R for every query q^i. One would then ask: if we can learn different percentage R^i for different queries, can the performance be further improved?

We define a feature vector $\mathbf{x}^i$ for each query q^i. For example, $\mathbf{x}^i$ can be q^i's low dimensional representation using the latent semantic analysis (LSA). A heuristic solution to the above question, for example, is to replace R with $R^i = \beta^T \mathbf{x}^i$, and optimize the objective function subject to β^T. However, the objective function defined this way becomes much more difficult to optimize, because we can no longer extract the common Δr_j shared by all queries.

This observation indicates that each R^i should be shared by at least a significant number of queries. To this end, we propose to leverage a decision tree structure. We assign every query q^i to one leaf node according to $\mathbf{x}^i$. For each leaf node t, we compute R_t for queries t using method discussed in Section 4.2. Since the goal of our problem is to find the optimal parameters instead of classification, we use the regression tree method (CART) [2].

[7]For example, there are four retrieved entities: 100 (position 1), 300 (position 2), 200 (position 3), 400 (position 4); the relevant entity is position 2. The relevant entity's percentage $z_i = \frac{3}{4}$.

Algorithm 1: PwCART(k)

Input: $\mathbf{x}^i, E^i, z^i, i = 1, 2, \cdots, n$
Output: Decision tree T where each leaf node represents a different percentage parameter

```
Grow a full tree according to splitting criterion 1 or 2;
if Splitting criterion = 2 then
    For each node t in the tree, use queries inside t to
    compute k range percentage parameter R_t;
Prune the tree with minimal cost-complexity pruning,
where the cost-complexity is defined as:
    C_{nα}(T) = C_n(T) + α|T̃|
Use the 0.5 SE rule to select the optimal tree;
```

When growing the tree, we first use the surrogate function C_n as the splitting criterion:

- **Splitting criterion 1**. Select dimension and separating value that minimizes C_n (Equation (3));

The original splitting criterion of CART is to minimize the mean square error (MSE). However, it is interesting to observe that this criterion is consistent with splitting criterion 1 to some extent. To explain this, imagine two different splits on the same data. Suppose that after the first split, data is perfectly separated into two small and isolating clusters; after the second split, however, data in the two clusters are still largely mixed. Between the two splits, the first one should have smaller variance. It should also have smaller C_n, because we can fit R to two very small regions. To this end, we define MSE as our second splitting criterion:

- **Splitting criterion 2**. Select dimension and separating value that minimizes the mean square error;

We summarize the steps of partition with CART (PwCART) in Algorithm 1. While splitting criterion 1 needs to reconstruct a new tree for every different k, criterion 2 only needs to build a single tree the whole time. Moreover, the time criterion 2 needs in constructing a single tree is also significantly less than criterion 1, because computing MSE is much faster than optimizing C_n. Overall, criterion 2 largely reduces the computation cost compared with criterion 1.

4.4 Time Cost and Rounding

Testing complexty. Upon receiving query q^i, our first method using DP (Section 4.1) takes $O(k|E^i|^2 + |E^i|^3 \log |E^i|)$ time to generate k ranges. Our second method (both Section 4.2 and Section 4.3) takes constant time to generate the percentage R^i, but the R^i still needs to be converted to its corresponding absolute values P^i. There are two ways to convert the percentages to absolute values: first, by sorting every $e \in E^i$ by $v(e)$, which takes $O(|E^i| \log |E^i|)$; second, by applying the k-th smallest element algorithm, e.g., quickselect[8], which takes $O(k|E^i|)$. When $|E^i|$ is large, this step can also be time consuming. However, we always need to scan E^i for at least one time. This is because after P^i is generated, we need to map each $e \in E^i$ to one of the $[p_{j-1}, p_j)$'s to obtain refined sub-lists. Therefore, the second method does not increase the time complexity with respect to $|E^i|$.

[8]https://en.wikipedia.org/wiki/Quickselect

Rounding. To generate the final ranges, we may need to round the floating numbers in P^i generated by our algorithms. Precision kept by this rounding depends on the specific application scenario. For price of products, users may be expecting more friendly designs, thus they may prefer "Below 100, Above 100" to "Below 99.7, Above 99.7". In other applications such as distance between locations, users may be more acceptable to higher precision such as "Below 11.7 miles, Above 11.7 miles". The rounding precision can be tuned as a parameter for trade-off between interpretability and ARR.

5. EXPERIMENTS

In this section, we conduct experiments on both our proposed methods to answer the question we raised in Section 1, i.e., can we leverage previous search logs to improve the results on test collection?

5.1 Dataset

Evaluation of the proposed models is challenging since no previous work has studied our problem setting described in Section 3. Therefore, we have to construct our own dataset.

We collect a two-month search log from walmart.com between 2015/10/22 and 2015/12/22. Since the size of the entire log is large and intractable on a single machine, we only keep the data from two categories: 'Laptop' and 'TV', because they are two of the most essential categories on the website with significant amount of traffic. Our data contains multiple numerical facets, e.g., *screen size*. We select the *price* facet for experiment, because it is the most important numerical facet with the largest amount of user engagement in our dataset.

For each category, we separate the earlier 70% as training data and latter 30% as testing data (according to timestamps). After the separation, Laptop contains 2,279 training queries and 491 testing queries, TV contains 4,026 training queries and 856 testing queries.

5.2 Experimental Results

In this section, we conduct an experimental study on the four methods we discuss in Section 4: `quantile`: simple method that partitions retrieved entities into equal-sized refined lists; `dp`: our first proposed method (Section 4.1) that optimizes expected ARR with dynamic programming; `powell`: our second proposed method (Section 4.2) where we use Powell's method for non-smooth optimization; `tree`: our decision tree method (Section 4.3).

5.2.1 Overall Comparative Study

Table 2 shows the ARR of the four methods. For every method we report the best tuned ARR by varying its parameters. We can see that the overall performance of `tree` is the best among all; `powell` and `dp` are next, with `powell` slightly better in Laptop and `dp` slightly better in TV; `quantile` has the worst performance in Laptop, and similar performance as `powell` in TV. On the other hand, if we vertically compare the performance of Laptop and TV, we can see that the ARR of `quantile` and `dp` is slightly better in TV, while that of `powell` and `tree` is better in Laptop.

To test the statistical significance in these results, we run the T-test[9] between `quantile`, `dp` and `tree`. We skip the T-test on `powell` because it uses the same methodology as `tree`, and the latter one is more general. The T-test results show that `tree` significantly ourperforms `quantile` and `dp` on the majority of experiments in Laptop category; the only exception is `tree` vs. `dp` when $k = 2$. This may be because when $k = 2$, the parameter space is the smallest among all, so the flexibility of parameterized method is reduced. On the other hand, the difference between `dp` and `quantile` method is not significant, and in TV category, none of the methods is significantly better than any other method.

Following the above comparison, we can observe one phenomenon: the parameterized methods `tree` and `powell` has especially better performance on Laptop data than TV. How does this happen?

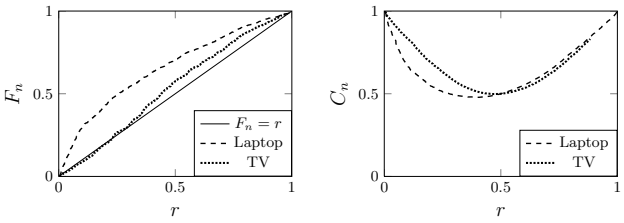

Figure 2: F_n and C_n for Laptop and TV when $k = 2$

	$k = 2$	$k = 3$	$k = 4$
`exhaustive`	31.72	21.27	16.14
`quantile`	31.85	21.30	16.19

Table 3: Optimal ARR vs. `quantile`'s ARR for 'TV'

To answer the above question, recall that the parametrized methods optimizes a surrogate function C_n, which is based on the empirical CDF F_n of relevant entity e^i's percentage z^i (Equation 3). In Figure 2 we plot the values of C_n and F_n when $k = 2$. These plots provide some insights for answering the above question: compared with Laptop, the empirical CDF of TV is very close to $F_n(r) = r$. If we plug $F_n(r) = r$ into Equation 3, the resulting C_n is a quadratic function, where the minimum is at $r = 0.5$. Consequently, observe that the minimum of TV's C_n in the right plot is very close to $r = 0.5$. Since $r = 0.5$ is exactly the quantile method when $k = 2$, this means when the empirical CDF is close to linear, it is difficult for the parameterized method to achieve a lower $C_n(R)$ than the quantile method.

However, our actual evaluation metric is ARR instead of $C_n(R)$. What about their ARR values? In Table 3 we show the actual values of the optimal ARR for TV by exhaustively search through the entire space of size $\binom{|E^i|-1}{k-1}$. When $k > 4$, this exhaustive search is intractable, therefore we only experiment on $k \leq 4$. We also display the ARR by `quantile` method for comparison. We can observe from Table 3 that the ARR of quantile method is indeed also very close to optimal. Therefore, we have explained the reason why parameterized methods has worse performance on the TV data.

5.2.2 Comparative Study on Non-smooth Optimization Methods

In this section we conduct comparative study on the performance of different non-smooth optimization methods. We study five non-smooth algorithms. Besides the aforementioned 1) `powell` and 2) `nelder-mead` methods, we also study: 3) `cg`: conjugate gradient method in non-smooth case; 4) `bfgs`: second order optimization method in non-smooth case; and 5) `slsqp`: sequential least square programming. For all

[9]https://docs.scipy.org/doc/scipy-0.15.1/reference/generated/scipy.stats.ttest_ind.html

		quant.	dp	powell	tree	tree vs. dp		tree vs. quant.		dp vs. quant.	
						p	t	p	t	p	t
Laptop	$k=2$	33.27	30.15	31.63	**28.00**	0.32	-0.98	9e-3	-1.45	0.15	-1.45
	$k=3$	22.07	21.22	19.95	**17.62**	0.03	-2.18	5e-4	-3.50	0.61	-0.50
	$k=4$	16.76	16.47	15.28	**13.29**	0.02	-2.23	3e-4	-3.63	0.83	-0.20
	$k=5$	13.55	13.43	11.94	**10.72**	0.04	-2.05	3e-4	-3.65	0.92	-0.09
	$k=6$	11.33	11.03	10.15	**9.03**	0.04	-2.02	2e-4	-3.69	0.76	-0.29
TV	$k=2$	31.85	30.99	31.73	**30.78**	0.89	-0.12	0.49	-0.68	0.60	-0.52
	$k=3$	21.30	20.88	21.43	**20.75**	0.89	-0.12	0.60	-0.51	0.69	-0.38
	$k=4$	16.19	15.95	16.30	**15.57**	0.63	-0.47	0.43	-0.78	0.76	-0.29
	$k=5$	13.08	12.83	13.18	**12.62**	0.75	-0.31	0.47	-0.72	0.70	-0.37
	$k=6$	10.95	10.64	10.98	**10.48**	0.76	-0.30	0.37	-0.89	0.57	-0.55

Table 2: Comparative study on the ARR of four methods. The ARR metric can be interpreted in this way: when the number of partitioned ranges is 6, users needs to read 11.33 entities in average with `quantile` method; while she only needs to read 9.03 entities in average with `tree` method. `dp`, `powell` and `tree` uses the same amount of training data for fair comparison.

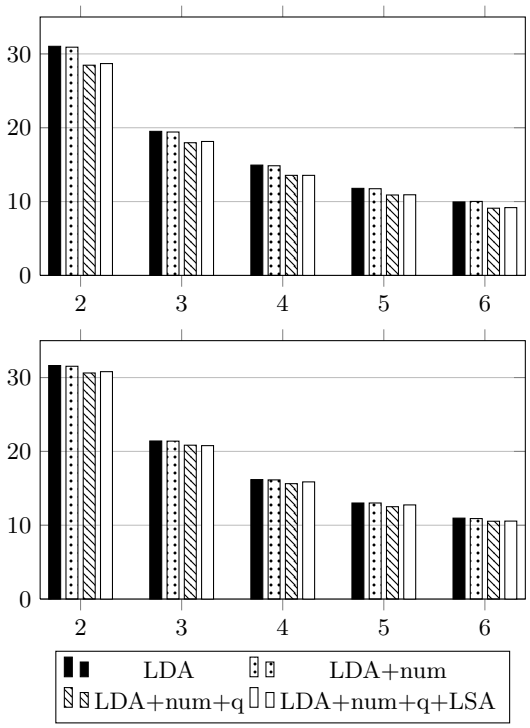

Figure 3: Compare importance of different feature groups: ARR for $k = 2, \cdots, 6$. Above: Laptop; below: TV

the five methods we use the implementation in Python library[10]. For each algorithm, we run 5-folds cross validation to tune the error tolerance as well as to find a good starting point. We repeat this process for 50 times and report the ARR and running time of each setting in Table 4. Due to the space limit, we only report the average values over $k = 2, \cdots, 6$.

From Table 4 we can see that the five algorithms have slightly different performances: `slsqp` has the best performance in Laptop and `powell` has the best performance in TV. `powell` and `nelder-mead` has the largest time cost, while `bfgs` is the fasted algorithm among all. This can be explained by the fact that `bfgs` is a second order method, while `Powell` and `nelder-mead` does not leverage the gradient information compared with the other three.

[10]https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.minimize.html

		powell	bfgs	nelder	cg	slsqp
avg ARR	L	17.77	17.58	17.78	17.60	**17.50**
	T	**18.70**	18.76	18.74	19.06	18.76
time	L	0.024	**0.007**	0.028	0.012	0.027
	T	0.022	**0.008**	0.026	0.009	0.009

Table 4: Compare different non-smooth optimization methods: average ARR and running time over $k = 2, \cdots, 6$.

5.2.3 *Comparative Study on Decision Tree Features*

Since decision tree method (Section 4.3) uses feature $\mathbf{x}^i$ for each query q^i, in this section, we study the importance of different feature groups in the final ARR. We use three groups of features for $\mathbf{x}^i$:

Semantic representation for q^i: we use both latent semantic analysis (LSA) and latent Dirichlet allocation (LDA). For each method the dimension is set to 20. In order to leverage more co-occurences, both method are run on the mixture of three categories instead of running on each category separately;

Number of explicitly mentioned facets in q^i: we use Stanford Named Entity Recognizer (NER) [11] to label the explicitly mentioned facets in each query. For example, for query "17 in refurbished laptop", the two explicitly mentioned facets are: screen size=17 and condition=refurbished. We manually label 40% of the queries for training, and the rest are output by the recognizer. The intuition behind leveraging this feature is that when user mentions more facets, it is more likely she is looking for higher profiled entities;

Absolute values of numerical facets in E^i: we use the quartile facets, which are the absolute values of the 25%, 50% and 75%th smallest facets in E^i. The intuition behind leveraging this feature is that when the retrieved entities are all very expensive, user may prefer relatively less expensive entities in the list;

The ARR of four combinations of feature groups are shown in Figure 3: (1) LDA (dimension=20): first, we only use the 20 features from LDA; (2) LDA + num (dimension=21): second, we add the number of explicitly mentioned facets in q^i; (3) LDA + num + q (dimension=24): third, we add the 3 quartile statistics; (4) LDA + num + q + LSA (dimen-

[11]http://nlp.stanford.edu/software/CRF-NER.shtml

sion=44): finally, we add the 20 features from LSA. In this experiment the splitting criterion of decision tree is fixed to criterion 2 and non-smooth optimization method is fixed to Powell's method.

Figure 3 shows that among all groups of features, the quartile features most significantly reduce the ARR. The number of explicitly mentioned facets does not significantly improve the ARR over LDA features. Given that LDA features already exist, adding the LSA features also do not significantly improve the ARR, and actually hurts the ARR in many cases.

5.2.4 Comparative Study on Decision Tree Splitting Criterion

In Section 4.3, we discuss the usage of two splitting criteria for building the decision tree. Recall the first criterion is to minimize $C_n(R)$ (Equation 3), while the second criterion is to minimize the mean square error. Therefore, we denote the first criterion as `nonsquare` and the second criterion as `square`. In this section, we study the influence of splitting criterion on the performance of decision tree. In order to make a comprehensive comparison, we look into three trees under each criterion: first, the fully grown tree without pruning, denoted as full; second, the smallest tree after pruning, which only contains the root node and two leaf nodes, denoted as min; third, the best ARR among all the pruned trees and the fully grown tree, denoted as best. In this experiment $\mathbf{x}^i$ is fixed to LDA + num + q and non-smooth optimization method is fixed to Powell's method.

In Figure 4 we show p values in the T-test results between the two criteria. When criterion 2 is better, we plot the p value in positive (`square`); when criterion 1 is better, we plot the p value in negative (`nonsquare`).

From Figure 4 we can see that the difference between the two criteria are basically consistent over $k = 2, \cdots, 6$. Although none of the p values is small enough to show statistical significance, we can still observe a few phenomena: first, the best ARR of `nonsquare` is slightly better than `square`; second, the smallest tree with two leaves of `nonsquare` is more significantly better than `square`; third, when the tree is fully grown, the ARR of `square` is instead better.

The above observed phenomena can be naturally explained: since the splitting criterion of `nonsquare` is to optimize C_n which approximates ARR, it is expected to achieve better ARR than `square`, for the same reason its smallest tree should also have better performance. Meanwhile, due to the scarcity of data samples in leaf nodes, the fully grown tree of `nonsquare` should be more overfitting than `square`, because it tries to fit ARR in every possible step.

5.2.5 Comparative Study on $p(e)$

	$k=2$	$k=3$	$k=4$	$k=5$	$k=6$
L	63.44	59.65	55.98	54.78	51.75
T	61.78	60.42	59.39	58.29	57.16

Table 5: ARR using $p(e) \propto 1/rank(e)$

In Section 4.1 we estimate the expected ARR based on the relevance model $p(e)$. In this section we study the performance of the DP algorithm using different $p(e)$'s. First, we study a combination of the query relevance and the category

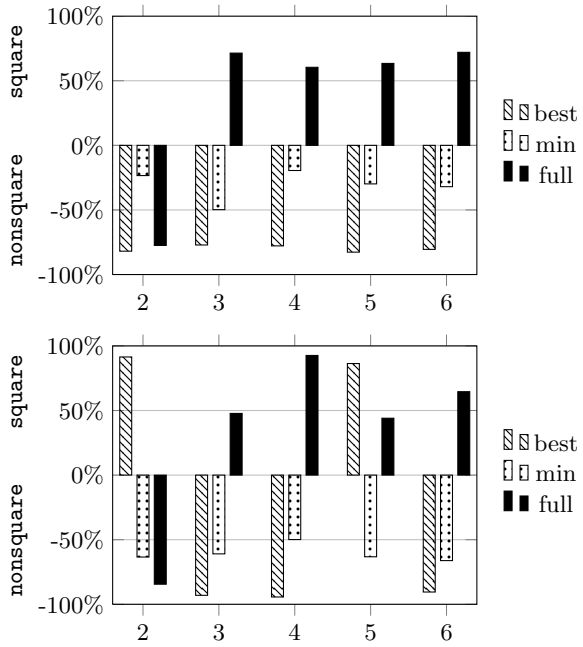

Figure 4: Compare different splitting criteria for decision tree method: p-value in T-test between minimizing mean square error (`square`) and minimizing C_n (`nonsquare`). Above: Laptop; below: TV

relevance models:

$$\begin{aligned} p(e) &= \lambda p_q(e) + (1-\lambda) p_{cate}(e) \\ p_{cate}(e) &\propto \#click(e, cate) \\ p_q(e) &\propto \#click(e, q) \end{aligned}$$

where $p_{cate}(e)$ is the number of clicks in the search log on entity e under category $cate$; $p_q(e)$ is the number of clicks on e under query q. The search log used to estimate this model is the training data described in Section 5.1. Therefore this method uses the same amount of training data as the parameterized methods. We use the best tuned parameter $\lambda = 0.5$ for ARR. The results using this relevance model is shown in Table 2.

Second, we study when $p(e)$ is estimated based on their original ranking on walmart.com, i.e., $p(e) \propto 1/rank(e)$. The ARRs are shown in Table 5. This method does not utilize the training data, and we can see its results are significantly worse than when using training data.

6. CONCLUSION

In this paper, we introduced a new problem of numerical facet range partition for faceted browsing. We proposed a novel evaluation metric ARR to quantify the quality of a partition based on the browsing cost for user to navigate into relevant entities. We proposed two methods that leverage search logs to improve the optimality of resulting partition. Evaluation results on an industry data show that our methods can significantly outperform the quantile baseline method. Among them, the second approach based on parameterization and learning is more robust and efficient, therefore can be directly applied to any search engine that supports numerical facet browsing. Our work opens up an interesting new area of study with many promising future directions. For example, we can more realistically simulate the effect of browsing by modelling the different browsing costs of different entities, and examine user behavior assumptions

that goes beyond sequentially scanning and stopping at the relevant entity. It is also a possible direction to explore how to partition the categorical facets. Finally, it would be interesting to consider how to uniformly model the complicated interactive decision process in faceted browsing system by simultaneously considering ranking and browsing.

7. APPENDIX

In this section, we give the proofs for two theorems which provide some useful insights on the convergence rate and sample complexity of the learning objective function in our second method (Equation (3)). Both theorems leverages the Dvoretzky-Kiefer-Wolfowitz (DKW) inequality [4] and the following property:

PROPERTY 1. *For any real value sequence $x_1, \cdots, x_n$ and $y_1, \cdots, y_n$:*

$$|\sum_{l=1}^{m} x_l y_l| \leq \sum_{l=1}^{m} |x_l| \times \max_{l'} |y_{l'}|$$

THEOREM 1. *Given query number n, suppose the relevant percentages $z^1, \cdots, z^n$ defined in Section 3 are independent and identically distributed. Let F_n denote their empirical CDF:*

$$F_n(r) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}[z^i < r]$$

and F denote the true CDF. Suppose C_n is defined by Equation (3), and C is the true function of C_n:

$$C(R) = \sum_{j=1}^{k} \Delta r_j \times (F(r_j) - F(r_{j-1}))$$

If the number of ranges is set to k, we can prove that given a constant $\epsilon > 0$:

$$\mathbb{P}[\sup_R |C_n(R) - C(R)| > \epsilon] \leq 2e^{-2n\epsilon^2/(k-1)^2}$$

PROOF. *Using Property 3:*

$$\begin{aligned} |C_n(R) - C(R)| &= |\sum_{j=0}^{k} (F_n(r_j) - F(r_j)) \times (\Delta r_j - \Delta r_{j+1})| \\ &\leq \sum_{j=0}^{k} |F_n(r_j) - F(r_j)| \\ &= \sum_{j=1}^{k-1} |F_n(r_j) - F(r_j)| \end{aligned}$$

where $\Delta r_0 = \Delta r_{k+1} = 0$. If $\sup_R |C_n(R) - C(R)| > \epsilon$, denote $\arg\max_R |C_n(R) - C(R)|$ as $R^0 = [r_1^0, \cdots, r_{k-1}^0]$, therefore:

$$\sum_{j=1}^{k-1} |F_n(r_j^0) - F(r_j^0)| > \epsilon$$

For at least one these j's we must have $|F_n(r_j^0) - F(r_j^0)| > \frac{\epsilon}{k-1}$ and thus $\sup_{r \in (0,1)} |F_n(r) - F(r)| > \frac{\epsilon}{k-1}$, therefore

$$\begin{aligned} & \mathbb{P}[\sup_R |C_n(R) - C(R)| > \epsilon] \\ \leq\ & \mathbb{P}[\sup_{r \in (0,1)} |F_n(r) - F(r)| > \frac{\epsilon}{k-1}] \\ \overset{DKW}{\leq}\ & 2e^{-2n\epsilon^2/(k-1)^2} \quad \square \end{aligned} \tag{4}$$

Theorem 1 describes the convergence rate of $C_n(R)$ for any point in the simplex space Δ^k. As k increases, bound (4) becomes looser. However, under certain setting, this bound will not increase with k. We can show it with Theorem 2 and Theorem 3 below:

THEOREM 2. *Suppose we have the same setting as Theorem 1, but in addition, the true CDF F is strongly concave. Denote $\arg\min_R C(R)$ as $R^* = [r_1^*, \cdots, r_{k-1}^*]$, then the widths of R^* is monotonously non-decreasing:*

$$\Delta r_1^* \leq \Delta r_2^* \leq \cdots \leq \Delta r_k^* \tag{5}$$

PROOF. *Since F is strongly concave, for any R and any pair of adjacent ranges $[r_j, r_{j+1})$ and $[r_{j+1}, r_{j+2})$ in R, we have:*

$$\frac{\Delta F(r_{j+2})}{\Delta r_{j+2}} < \frac{\Delta F(r_{j+1})}{\Delta r_{j+1}} \tag{6}$$

where $\Delta F(r_{j+1}) = F(r_{j+1}) - F(r_j)$.

Given the optimal point R^*, now consider R', which is same as R^* except for replacing r_{j+1}^* with $(r_j^* + r_{j+2}^*)/2$. Since R^* is the optimal point:

$$C(R^*) \leq C(R') \tag{7}$$

By canceling the same terms on the L.H.S. and R.H.S. of (7) we can get:

$$\begin{aligned} & \Delta r_{j+1}^* \Delta F(r_{j+1}^*) + \Delta r_{j+2}^* \Delta F(r_{j+2}^*) \\ \leq\ & \tfrac{\Delta r_{j+1}^* + \Delta r_{j+2}^*}{2} (\Delta F(r_{j+2}^*) + \Delta F(r_{j+1}^*)) \\ \Rightarrow\ & (\Delta r_{j+1}^* - \Delta r_{j+2}^*)(\Delta F(r_{j+2}^*) - \Delta F(r_{j+1}^*)) \geq 0 \end{aligned} \tag{8}$$

Suppose (5) is not true, and there exists a j such that $\Delta r_{j+2}^* < \Delta r_{j+1}^*$. It follows from (9) that $\Delta F(r_{j+2}^*) < F(r_{j+1}^*)$, therefore $(\Delta r_{j+1}^* - \Delta r_{j+2}^*)(\Delta F(r_{j+2}^*) - \Delta F(r_{j+1}^*)) < 0$, which contradicts with (8). $\square$

THEOREM 3. *Suppose we have the same setting as Theorem 1. In addition, the true CDF F is strongly concave and $R^* = \arg\max_R C(R)$. Denote $\mathcal{R}^*$ as a small enough region near R^* where (5) stays true, then for constant $\epsilon > 0$:*

$$\mathbb{P}[\sup_{R \in \mathcal{R}^*} |C_n(R) - C(R)| > \epsilon] \leq 2e^{-2n\epsilon^2} \tag{9}$$

PROOF. *Following (5), for any $[r_j, r_{j+1})$ in $R \in \mathcal{R}^*$, $|\Delta r_j - \Delta r_{j+1}| = \Delta r_{j+1} - \Delta r_j$. Using Property 3:*

$$\begin{aligned} & \sup_R |C_n(R) - C(R)| \\ \leq\ & \sup_R (\Delta r_k - \Delta r_1) \times \max_j |F_n(r_j) - F(r_j)| \\ \leq\ & \sup_{r \in (0,1)} |F_n(r) - F(r)| \end{aligned}$$

Following DKW inequality we get (9). $\square$

We may combine the results in Theorem (1-3) with experimental results in Section 5 and draw some conclusions. Recall that our second method achieves better experimental results on 'Laptop' category than 'TV'. From Figure 2, we can observe that the true CDF of 'Laptop' is strongly concave while that of 'TV' is mostly linear. Meanwhile, Theorem 3 shows that when the true CDF is strongly concave, $C_n(R)$ also has better convergence rates. The two results demonstrate some consistency between theoretical analysis and experimental results.

8. REFERENCES

[1] J. Acharya, I. Diakonikolas, C. Hegde, J. Z. Li, and L. Schmidt. Fast and near-optimal algorithms for approximating distributions by histograms. In T. Milo and D. Calvanese, editors, *PODS*, pages 249–263. ACM, 2015.

[2] L. Breiman, J. Friedman, R. Olshen, and C. Stone. *Classification and Regression Trees*. Pacific Grove, 1984.

[3] R. Brent. *Algorithms for minimization without derivatives*. Prentice-Hall, 1973.

[4] A. Dvoretzky, J. Kiefer, and J. Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. The Annals of Mathematical Statistics, pages 1397–1400. ACM, 1956.

[5] M. A. Hearst. *Search User Interfaces*. Cambridge University Press, 1 edition, 2009.

[6] H. V. Jagadish, N. Koudas, S. Muthukrishnan, V. Poosala, K. C. Sevcik, and T. Suel. Optimal histograms with quality guarantees. In A. Gupta, O. Shmueli, and J. Widom, editors, *VLDB*, pages 275–286. Morgan Kaufmann, 1998.

[7] A. Kashyap, V. Hristidis, and M. Petropoulos. Facetor: cost-driven exploration of faceted query results. In J. Huang, N. Koudas, G. J. F. Jones, X. Wu, K. Collins-Thompson, and A. An, editors, *CIKM*, pages 719–728. ACM, 2010.

[8] S. Liberman and R. Lempel. Approximately optimal facet selection. In *Proceedings of the 27th Annual ACM Symposium on Applied Computing*, SAC '12, pages 702–708, New York, NY, USA, 2012. ACM.

[9] M. Muralikrishna and D. J. DeWitt. Equi-depth histograms for estimating selectivity factors for multi-dimensional queries. In H. Boral and P.-Ã. Larson, editors, *SIGMOD Conference*, pages 28–36. ACM Press, 1988.

[10] J. A. Nelder and R. Mead. A simplex method for function minimization. *Computer Journal*, 7:308–313, 1965.

[11] S. B. Roy, H. Wang, G. Das, U. Nambiar, and M. K. Mohania. Minimum-effort driven dynamic faceted search in structured databases. In J. G. Shanahan, S. Amer-Yahia, I. Manolescu, Y. Zhang, D. A. Evans, A. Kolcz, K.-S. Choi, and A. Chowdhury, editors, *CIKM*, pages 13–22. ACM, 2008.

[12] K. Sparck Jones and P. Willett, editors. *Readings in Information Retrieval*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1997.

[13] R. Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

[14] R. van Zwol, B. SigurbjÃűrnsson, R. Adapala, L. G. Pueyo, A. Katiyar, K. Kurapati, M. Muralidharan, S. Muthu, V. Murdock, P. Ng, A. Ramani, A. Sahai, S. T. Sathish, H. Vasudev, and U. Vuyyuru. Faceted exploration of image search results. In M. Rappa, P. Jones, J. Freire, and S. Chakrabarti, editors, *WWW*, pages 961–970. ACM, 2010.

[15] B. Wu, P. A. Szekely, and C. A. Knoblock. Minimizing user effort in transforming data by example. In T. Kuflik, O. Stock, J. Y. Chai, and A. KrÃijger, editors, *IUI*, pages 317–322. ACM, 2014.

[16] Y. Yue, T. Finley, F. Radlinski, and T. Joachims. A support vector method for optimizing average precision. In W. Kraaij, A. P. de Vries, C. L. A. Clarke, N. Fuhr, and N. Kando, editors, *SIGIR*, pages 271–278. ACM, 2007.

[17] C. Zhai and J. D. Lafferty. A study of smoothing methods for language models applied to ad hoc information retrieval. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2001*, pages 334–342, 2001.